

Amendments to the Specification:

Amendments to the specification below are keyed to the Paragraph numbering in the corresponding Publication No. US-2003-0103521-A1.

[0005] The commonly used ISO OSI communications model specifies a seven layer communications protocol stack comprising a physical layer, link or media access control (MAC) layer, network ~~layer~~ layer, transport ~~layer~~ layer, session layer, presentation layer and an application layer. The main functionality of the MAC layer and associated protocol, is to provide an effective method of accessing the physical channel, which in this case is the power line. Along with basic channel access, the MAC protocol ideally also provides (1) efficient and reliable host packet and session transport for the different kinds of possible connections, including unicast, multicast and broadcast, (2) efficient addressing scheme for the stations connected to the network, (3) repeater functionality for large networks with stations spread over large distances and (4) quality of service functionality.

[0010] Another characteristic of power line based networks is that stations on different logical networks may share the same physical network. In many cases, a home, enterprise or other premise includes more than one communication network. Each communication network may be made up of a plurality of nodes with each network comprising at least two nodes. All nodes of the same network implement the same communication technique and are able to ~~communication~~ communicate with each other thus permitting interoperability (assuming that the propagation conditions over the media enable communication). Nodes from different networks may implement different ~~communications~~ communication techniques, in which case they are not able to communicate with each other. In addition, the propagation characteristics of the shared media (e.g., the powerline grid) may have large variations and irregularities. This results in large variations in the attenuation over the communication path between two given nodes.

[0011] Since the powerline grids of neighboring residences are physically connected via the power distribution network, the common media of Figure 1 might refer to the powerline grid

of a single residence or to the powerline grids of several neighboring residences (e.g., several apartments in a building).

[0012] Thus, the entire physical network is vast, possibly comprising a large number of houses or enterprises, and cannot be separated into small logical cells, e.g., a single house. This is because ~~often~~, neighborhoods are often wired such that many homes or enterprises in the same area are connected to the same electrical phase of the utility wiring.

[0013] Another characteristic of the power line network is that channel attenuation ~~maybe~~ may be non-symmetric, i.e. the attenuation may be different for opposite directions of signal propagation. In other words, the impedance as seen by the stations on opposite sides of transmission is different. This may cause power mismatching or other problems during transmission and reception.

[0014] Yet another characteristic of power line networks is that they are ad hoc in nature meaning that stations may be added to or removed from the network at any time. In such types of networks, there is no system administrator or installer as there are in other well controlled networks such as telecommunication carrier WANs, corporate networks, etc. No calibration measurements are taken and each end user buys and installs devices independent of and without the knowledge of other users of the network.

[0021] If the physical network is very large, however, containing many logical networks, simple flooding will cause the original transmission to be unnecessarily transmitted an extremely high number of times causing congestion ~~[[in]]~~ throughout the entire physical network. Reducing the number of these transmissions is possible by predefining a number of logical networks and having only stations in the same logical network repeat the original transmission. The drawback ~~[[in]]~~ of this scheme is that different logical networks must be defined in the MAC layer ~~that~~ which in some applications is an impractical requirement.

[0106] When a source station detects a link failure while communicating or attempting to communicate with its destination, it can utilize an intermediate station to transfer frames to the

destination. The intermediate station functions as a second layer repeater[[,]]. In operation, the entire session, including retransmissions,[[[]]] is transmitted to the repeater station whereby the data frames are then transmitted to the destination station under a new session. The second layer repeater and repeater establishment processes are described in more ~~details~~ detail infra in Section 4.

[0109] The MAC protocol also provides a mechanism for network synchronization. The network synchronization scheme provides a synchronization of better than 8 μ sec between neighboring stations in the network. The synchronization scheme is described in more detail infra in Section 5 and in ~~U.S. Application Serial No. 09/822,939, filed March 30, 2001~~ U.S. Patent No. 6,907,472, entitled "Distributed Synchronization Mechanism For Shared Communications Media Based Networks," similarly assigned and incorporated herein by reference in its entirety.

[0127] The host specifies a multicast transmission by one of the following two ways: (1) using an Ethernet multicast MAC addresses (0x01 : 00 : 5E : xx : xx : xx) or (2) using a predefined multicast group defining a plurality (e.g., 16) of stations within the group. Ethernet multicast MAC addresses are translated into broadcast transmissions and are transmitted as described above in Section 2.1.2. [[]]

[0132] To support higher layer synchronization, the MAC protocol of the present invention is adapted to enable both the transmitting and receiving station to synchronize to the clock of the transmitting MAC. To accomplish this, when requested, the transmitter MAC adds a time stamp t_1 to the packet. The receiver MAC adds its time stamp to the packet before sending it to the transmitter station. Both the transmitter station and receiver station have access to the clock of their respective MAC layers. The receiver host calculates the time at the transmitter MAC using the following:

[0139] In the MAC protocol, communications between stations take place during sessions. A session is defined as the logically related transmissions between two stations including contention frames (RTS/CTS), data frames, ACK frames and any subsequent retries.

Packets received from the host are associated together so as to form sessions. The association process allows for the segmentation of a large packet into [[a]] multiple sessions or the accumulation of several small packets into one session.

[0140] The two types of ~~sessions-transported~~ session transport supported include long session transport (LST) and short session transport (SST). In LST sessions, the channel is captured and reserved using the RTS/CTS mechanism. The session is composed of up to MaxFragNumber variable size frames and includes any required retransmissions. In SST sessions, the channel is not reserved prior to session transport. The session is composed of a single frame, and includes any required retransmissions. The size of the fragment is limited to FragmentThreshold bytes. Detailed descriptions of the two session transport processes are given in the following sections.

[0142] An LST session may comprise up to MaxFragNumber data frames, whereby the size of each frame is determined using the process detailed below. A diagram illustrating the long session transport (LST) transmission process of the MAC protocol of the present invention is shown ~~in~~ in Figure 6. Channel access is accomplished using an RTS control frame after the contention interframe space (CIFS) and backoff interval. A response interframe space (RIFS) separates the transmitted packets. Following the CTS packet, the data frames are transmitted followed by the ACK and any retransmissions.

[0177] A station may also transmit a media release message when it determines the channel can be released. This may be determined when the virtual carrier sense (VCS) signal indicates that the communications channel is busy (i.e. via a reservation calculation) but the channel is ~~not~~ no longer required or can no longer be used. The media release message comprises an updated value for the reservation counter. The station sending the media release message may (1) increase the reservation counter value if it is found the current reservation time was too short; (2) decrease the reservation counter value if it the current reservation time is too long; or (3) set the reservation counter value to zero in order to immediately release the channel if it is found the channel is no longer required or cannot be used.

[0193] In accordance with the present invention, medium sharing is achieved using a modified CSMA/CA mechanism with random backoff. The medium sharing implemented by the MAC protocol uses the following mechanisms: virtual carrier sense (VCS) and physical carrier sense (PCS), channel reservation, backoff and interframe space. Each mechanism is ~~described~~ described in the following sections.

[0206] A station decrements the backoff value only when the channel is not busy and during the contention windows of priorities equal to or lower ~~[[to]]~~ than the priority of the pending transmission. If transmission is detected during backoff countdown, the station defers transmission until the contention window in the next cycle having a priority equal to or lower than its priority and begins decrementing the backoff counter from its previous value.

[0209] The FCD signal is a signal passed from the physical layer to the MAC layer that indicates that a transmission may be starting. This signal arrives very quickly after the beginning of transmission but has a relatively high false alarm rate. The CD signal is also passed from the physical layer to the MAC layer that indicates with very high probability that a transmission is starting. This signal arrives a relatively long time after the start of transmission but has a very low false alarm rate.

[0212] The backoff algorithm of the present invention is intended to prevent the stations from competing for the channel at the same time. After the current transmission completes, a station may begin contending for the channel. The contention period is divided into time slots. Stations randomize how many time slots ~~[[it]]~~ must ~~defer~~ be deferred before trying to compete for the channel. If during the deferral time, another station begins transmission, the station waits for the next time the channel is free. The size of the time slot is determined such that stations that randomize a different deferral period will not collide on the channel. This means that the station that randomized a longer deferral period will detect that the second station has already begun transmission. The size of the time slot is usually set to the time it takes for stations to synchronize to the packet on the channel. The larger the time slot, however, the more overhead the backoff algorithm adds to the channel access mechanism.

[0214] While the channel is in a contention period (step 206), the station decrements the backoff time (step 210) after waiting CWMin of the higher priority (step 208). While the backoff counter is not zero (step 212), the method checks for the VCP and PCS to be idle and decrements the backoff counter. When the backoff counter reaches zero (step 212), transmission proceeds by attempting to reserve (or acquire) the channel as described hereinbelow in Section 6.5. Reservation of the channel is achieved by transmitting a RTS or RA frame (step 214). The protocol implements two special control packets designed to capture the channel: Request To Send (RTS) and Clear To Send (CTS). The RTS and CTS packets are used to capture and reserve the channel by both the source station and the receiving station.

[0245] The second collision situation is more complex since the hidden nodes will not receive acknowledgements and will wish to retransmit. This pattern is repeated until the number of retransmissions has been exhausted. The upper ~~communications~~ communication layers will be notified and as a result will generate the packet again resulting in an infinite loop until the upper communication layers close the connection.

[0253] All reservation times are compressed to fit the size of one byte before transmission. The byte is represented by the RES field. Each receiving station decompresses the RES field and updates the VCS timer using the decompressed reservation time. The reservation time is counted from the last received bit of the frame, which was also used to update the VCS timer. Any suitable compression and decompression technique can be used and is not critical to operation of the invention. The reservation time is calculated as described in Table 6 below. below.

[0245] ~~Table below.~~

[0256] A flow diagram illustrating the CTS method of the present invention is shown in Figure 24. The timeout is first set to the value of the CTS_Timeout (step 270). While a timeout has not occurred (step 272), the timeout is decremented (step 274) and a check is performed whether a CTS frame has been received (step 276). If [[so]] a CTS frame has been received, the

method returns CTS received (step 280), otherwise, the method continues with the timeout check in step 272. Upon timeout, the method returns with timeout (step 278).

[0266] Data is presented to the physical layer most significant byte first, most significant bit first wherein bit number 7 is the MSB of a byte. The diagrams presented herein ~~shown~~ shows the most significant bits or bytes to the left.

[0305] As described above, the 6 byte MAC address is used to identify each station on the network. The necessity of using 6 bytes for each station arises from the fact that a MAC address is a unique identification of a station not only in the logical network but also as a physical entity. In a logical network, however, logical ~~addresses~~ addresses can be assigned to reduce the size of the address space. For example, if the logical network has up to 256 stations, a one byte logical address is sufficient, thus reducing the address size by more than 80%. As described above, station addresses are used in the RTS and CTS frames, thus reducing the overhead of the addresses and thereby significantly reducing the overall overhead of the MAC layer.

[0306] The deployment of logical addresses in a centrally controlled network is performed as follows. Each station that joins the network is allocated a logical address by the central controller. In a distributed network, the process of assigning logical address must also be done in a distributed way. The main problem of assigning a logical address in a distributed network is making sure that a station does not choose a logical address that is already in use.

[0314] The link address assignment process can be accomplished using both long RTS and long CTS frames or by using two management frames: link address set (LAS) and link address set response (LASR). The LAS frame is sent by the initiating station and comprises the LDA portion of the LA. The LASR frame is sent by the destination station and comprises the LSA portion of the LA. The process of establishing link address is described in the following subsections.

[0320] The LA algorithm described above in combination with the large link address space (2^{16} possible link addresses) significantly ~~reduce~~ reduces the possibility of a collision. Duplicate link ~~address~~ addresses may, however, occur in some extreme cases. The MAC protocol of the present invention includes a mechanism to resolve duplicate link ~~addresses~~ addresses. The occurrence of a duplicate link address may cause either of the following effects: (1) the CTS response frames from two stations collide preventing the reception of either at the source station or (2) only one CTS response frame is received but from the wrong station, (3) both stations respond with CTS response frames, both of which are correctly received at the source station.

[0322] The second and third cases are likely to result in the packet being received by the wrong station. To resolve this type of duplicate address problem, the source station inserts the MAC address of the destination station into the packet. Upon receiving the packet, the destination station checks the MAC address, if there is no match, the destination discards the packet and established a new LA. Otherwise, processing proceeds as described above.

[0326] The reservation field in the ACK frame is calculated according to the transmission rate used by the transmitter. When the session is finished, the transmitter records the power level and rate used in the transmission. The power level and rate records are used ~~for~~ for the next transmission on the same link.

[0336] To adapt to the TCP protocol the initial PLR level is set to the lowest level, to ensure connectivity, and ~~the~~ IncPLRThreshold and DecPLRThreshold ~~begin~~ start from a small value, and increase with each PLR state, until they reach their final value.

[0337] In accordance with the invention, a station receives statistics from other stations using a request to ~~gathered~~ gather statistic frame (RGS). The stations that receive the request respond using a gathered statistics response (GSR) frame. The GSR frame is transmitted as a broadcast session with MAX_HTL initialized to zero. Optionally, each station is adapted to broadcast a GSR frame with MAX_HTL equal to zero periodically every StatInterval. In addition, each station broadcasts a management frame that includes the type of priority it

transmitted in the past StatInterval seconds. The stations add the received information to their statistics and calculate CWMin values accordingly.